

O

AR-010-106

DSTO-TR-0478

T

Evaluation of a Laboratory Automatic
Target Recognition system:
Air-to-Sea Image Sequence

Paul Miller, Peter Virgo,
Mike Royce and Steve Angeli

S

APPROVED FOR PUBLIC RELEASE

© Commonwealth of Australia

R

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

Evaluation of a Laboratory Automatic Target Recognition System: Air-to-Sea Image Sequence

Paul Miller, Peter Virgo, Mike Royce and Steve Angeli

**Land Space and Optoelectronics Division
Electronics and Surveillance Research Laboratory**

DSTO-TR-0478

ABSTRACT

The operation and evaluation of a laboratory automatic target recognition (ATR) system is described. The ATR system was evaluated against its ability to acquire and track a specific target in the presence of false targets and background clutter. A ten second air-to-sea image sequence that contained the infrared signatures of three ship targets and sea clutter was used as the test sequence. The evaluation involved assessing the system performance with default values for various system parameters. These were then varied and the effect on system performance measured. The ability to acquire and track the true target over a limited range of target orientations and in the presence of low clutter and false targets was demonstrated. Furthermore, it was found that acquisition was easier if the target was somewhat broadside rather than end-on to the sensor. A shortcoming of the system was that it was not robust to minimal target obscuration. In addition, its ability to discriminate between the true and false targets was somewhat limited, and resulted in the acquisition and tracking of false targets. The evaluation showed that the system discrimination was improved by optimisation of a number of parameters, however, this also resulted in a reduction of its ability to acquire and track the true target. In addition, those parameters that were most important in terms of their effect on the system performance were determined.

RELEASE LIMITATION

Approved for public release

DTIC QUALITY INSPECTED 2

DEPARTMENT OF DEFENCE

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

19970724 090

Published by

*DSTO Electronics and Surveillance Research Laboratory
PO Box 1500
Salisbury South Australia 5108*

*Telephone: (08) 259 5555
Fax: (08) 259 6567
© Commonwealth of Australia 1997
AR-010-106
February 1997*

APPROVED FOR PUBLIC RELEASE

Evaluation of a Laboratory Automatic Target Recognition System: Air-to-Sea Image Sequence

Executive Summary

Focal-area reconnaissance is the reconnaissance of an area of radius of one hundred kilometres. Focal-areas may be in the vicinity of an Australian asset, near an adversary's base, or in specific geographical areas. Such a mission is initiated, or cued, by the detection of a target following surveillance of a much broader area. A major requirement of the Australian Land Force's focal-area reconnaissance capability is an ability to acquire cued land vehicle targets in the Northern Australian environment. A technology that has the potential to enhance present and future systems that provide this capability is real-time automatic/aided target recognition (ATR). This potential arises from the ability of ATR systems to recognise targets in various types of sensor imagery. Under sponsorship from the Directorate Combat Force Development (Land), Land, Space and Optoelectronics Division is evaluating the potential of this technology for acquiring and tracking targets in electro-optic imagery.

A primary aim of the project has been to implement a concept demonstrator ATR system to get experience in the technology that can be translated into useful advice to the ADF. To achieve this a concept demonstrator has been developed that is based on the System for Passive Optical Target Recognition (SPOTR) developed by Lockheed Martin under the Transition of Optical Processors into Systems (TOPS) project in the USA. Both the SPOTR and our system employ optical processing hardware that results in reduced system power requirements, weight and size. Furthermore, the SPOTR was shown to be capable of good discrimination between different classes of land vehicles. As a result of the TOPS program ruggedised, mil-spec ATR systems are now on the verge of becoming commercially available. Therefore, the work reported herein provides a timely insight into these developments.

The task consists of a number of phases with associated reports. The present report details the experiments carried out on the concept demonstrator as originally constructed, along with results for a simple image sequence. At present the components of the concept demonstrator are being substantially upgraded: a subsequent report will describe the improved system, its performance on more demanding image sequences and present some recommendations on future directions.

The aim of the experiments reported here was to test the ATR system's ability to acquire and track a target in a single image sequence. Due to the extremely high clutter present in electro-optic imagery obtained from the land environment and the untested nature of the system, it was decided to test the system on a low-clutter infra-red image sequence of three ships at sea. Thus the results presented here are not intended to be hard quantitative assessments of the system performance; rather they

are an experimental record, a guide to the ongoing system upgrade and give qualitative insights into the likely performance of a real system.

Nevertheless, the report shows that the primary aim of the first stage of the task has been achieved; a working concept demonstrator has been constructed that provides useful insights into the strengths and limitations of the technology with respect to a specific image sequence. Specifically, the ability to acquire and track a target over a limited range of target orientations and in the presence of false targets was demonstrated. In addition, it was found that

- system performance was improved if the target was broadside to the sensor rather than end-on;
- the system lost target track if the target became partially obscured by a false target;
- there was a trade-off in the ability of the system to acquire and track the target and its ability to discriminate between the true target and false targets. High probability of acquisition resulted in the tracking of false targets in some instances, whilst low probability of acquisition resulted in the tracking of the true target only.

The last of these observations is partly due to the limited resolution of the system. The associated technology is advancing rapidly, however, and in the upgraded concept demonstrator speed and resolution will be increased by up to an order of magnitude. The next stage of the task will report on the performance of the upgraded system: nevertheless a significant amount of further work needs to be done particularly in training the system and applying it to more demanding image sequences from the land environment.

In conclusion, this technology concept shows reasonable promise for realising automatic target acquisition and tracking in specific applications by recognising the class of a target - it cannot identify a specific target. Furthermore, as a result of this evaluation the quality of the DSTO knowledge base in this technology area has improved significantly. Also, it has become apparent that future work should be performed within the context of a broader evaluation process that addresses the relevant technologies, such as ATR, from a system perspective.

Authors

Paul Miller

Land, Space and Optoelectronics Division

Paul Miller received his BSc and PhD in Physics from The Queen's University of Belfast in 1985 and 1989, respectively. He then gained a 15 month research fellowship from the Defence Research Agency in the UK, which was also based at Queen's. Since 1991 he has been a research scientist with the Land, Space and Optoelectronics Division of DSTO at Salisbury. His research interests include all aspects of optical correlators for automatic/aided target recognition and tracking. He has previously worked on the application of optical processors to real-time speckle metrology, and the design of procedures for the mapping of algorithms onto parallel processing array architectures.

Peter Virgo

Land, Space and Optoelectronics Division

Peter Virgo is a senior technical officer at DSTO with Associate Diplomas in Electronic Engineering, Computer Applications and Systems Analysis. His recent work includes software development on image processing hardware, analog and digital video image processing, image management, image acquisition system design and development and trials support. His work history includes electronic development and trials support for the Woomera Rocket Range, development of a Sonar Projector Buoy system, microprocessor based data acquisition and replay systems for a Laser Airborne Depth Sounder and the establishment of an Optical Signatures Measurement and Analysis Facility.

Mike Rocye

Land, Space and Optoelectronics Division

Michael Royce recieved his BEEE from the University of Adelaide in 1985. After joining DSTO in 1986 as a Professional Officer Class 1, he participated in a program to automatically identify targets in remotely sensed imagery using their multi-spectral signatures. At the beginning of 1989 he began work on a performance prediction project for the upgrade of the IRDS used in PC-3 Orion aircraft.

Following an 18 month working holiday in the UK, where he completed two software contracts for private companies, he returned to DSTO and began work in optical processing. In 1994 he was promoted to Professional Officer Class 2. Currently, he is responsible for implementing the digital hardware and software components of an automatic target recognition system.

Steve Angeli

Land, Space and Optoelectronics Division

Steven Angeli received his B.Sc.(Ma) from the University of Adelaide in 1984. In 1986 he commenced work as a Computer Systems Officer at DSTO. Since then he has worked in the areas of image processing and synthesis; the thermal behaviour of large masses; and filter generation for target recognition. He is currently employed as a Professional Officer in the Image Processing and Propagation Discipline of LSOD working on the statistical analysis of images.

Contents

1. INTRODUCTION.....	1
2. ATR TASK.....	2
3. OVERVIEW OF CONCEPT DEMONSTRATOR ATR SYSTEM.....	3
3.1 Description and functionality	4
3.1.1 Image Preprocessor.....	5
3.1.2 Filter Database.....	7
3.1.3 Correlator.....	9
3.1.4 Postprocessor.....	10
3.1.5 Image/Filter Management.....	12
3.2 Operation	15
4. SYSTEM EVALUATION.....	16
4.1 Diagnostic test	17
4.2 Relationship between diagnostic test result variation and laboratory temperature.....	17
4.3 Filter discrimination	19
4.4 System performance with default parameter values.....	20
4.5 Effect of varying the sdm parameter on system performance.....	23
4.6 Influence of S_{th} and T_{th} on system performance	26
4.7 Comparison of system performance with different search strategies.....	28
4.8 Comparison of system performance with different EPAR measurements.....	29
5. SYNOPSIS OF RESULTS	31
6. SUMMARY AND CONCLUSION.....	33
7. REFERENCES.....	34
APPENDIX 1	37
APPENDIX 2	37

1. Introduction

Focal-area reconnaissance is the reconnaissance of an area of radius one hundred kilometres in the vicinity of an Australian asset, near an adversary's base, or in specific geographical areas. Such a mission is initiated, or cued, by the detection of a target following surveillance of a much broader area. A major requirement of the Australian Land Force's focal-area reconnaissance capability is an ability to acquire cued land vehicle targets in the Northern Australian environment. One technology that may enhance present and future systems that provide this capability is real-time automatic/aided target recognition (ATR). ATR systems, once trained, have the ability to recognise targets in various types of sensor imagery. Through Image Processing and Propagation (IPP) Discipline, Land Operations Branch of the Land, Space and Optoelectronics Division (LSOD) is evaluating the potential of this technology for acquiring and tracking targets, by recognition, in electro-optic imagery.

A primary aim of the project has been to implement a concept demonstrator ATR system to get experience in the technology that can be translated into useful advice to the ADF. To evaluate the technology a concept demonstrator has been developed that employs an optical correlator - a class of analogue optical processor. This demonstrator is based on the System for Passive Optical Target Recognition (SPOTR) developed by Lockheed Martin under the Transition of Optical Processors into System (TOPS) project in the USA. The motivation for employing an optical correlator is that it results in reduced power requirements, weight and size of the ATR system. Furthermore, the correlation algorithm implemented by the optical correlator is a high quality, computationally intensive, classification algorithm. As a result of the TOPS program ruggedised, mil-spec optical correlator based ATR systems are now commercially available. Therefore, the work reported herein provides timely information in the light of these developments.

The task consists of a number of phases with associated reports. The present report details the experiments carried out on the concept demonstrator as originally constructed, along with their results for a simple image sequence. At present the components of the concept demonstrator are being substantially upgraded: a subsequent report will describe the improved system, its performance on more demanding image sequences and present some recommendations on future directions.

The approach adopted in these preliminary investigations was to test the ability of the optical correlator ATR system to acquire and track a target in a single image sequence. However, as extremely high clutter is present in electro-optic imagery obtained in the land environment, it was decided to first test the system on low clutter imagery obtained from the maritime environment. To evaluate the system the effect of changing various operating parameters on system performance was investigated.

The report begins with a description of the ATR task and test image sequence against which the system is to be evaluated. This is followed by a description of the concept demonstrator ATR system developed under this program. The experiments performed to evaluate the system are then described and the results presented in

section 4. This also includes a detailed analysis of the results. Section 5 summarises the major results from the experiments. The report ends with a summary and conclusion section that also includes a brief outline of future work to be performed.

To our knowledge this is the first time that an evaluation of this technology on electro-optic imagery from the Australian environment has been performed. The concept demonstrator developed under the project has provided useful insights into the strength and limitations of the technology. This knowledge will provide the basis for useful advice on the technology to major defence projects.

2. ATR Task

For this evaluation the ATR system was tasked to acquire and track a single target in an image sequence obtained in a low clutter maritime environment. To facilitate the testing a suitable sequence of test imagery was chosen from the LSOD Image Database Facility. The sequence was obtained during the ARMADA '86 field trial [2] from an infrared sensor mounted on an airborne platform (C47 aircraft), and subsequently recorded on video tape.

The test image sequence features three ship targets sailing in convoy. The sequence is of approximately ten seconds duration and consists of 288 image frames. Each image was grabbed from video tape, digitised and stored on the hard disk of a host PC. The digitised images contained 512×512 8-bit pixels. Figure 1 shows the first raw image frame of the sequence. (For convenience the upper left ship will be called ship 1 henceforth, the middle ship, ship 2, and the lower right ship, ship 3.) As illustrated by Figure 1, the sequence begins with the plane to the rear of, and to one side of, the convoy. During the course of the sequence the plane moves round the rear of the convoy until it is on the other side of the convoy. For these investigations the ATR system was tasked with acquiring and tracking ship 2 during the course of the image sequence.




Figure 1: First image frame in test sequence

3. Overview of concept demonstrator ATR system

The ATR system used in the evaluation is based on the system developed for the TOPS program. As such, the main feature of both systems is that they employ an optical correlator. However, although they implement the same correlation algorithm, the correlator used in these investigations is a laboratory system whose performance is limited compared to the correlator used in the TOPS program [3] or current state-of-the-art correlators that are commercially available (Appendix 1). In particular, the laboratory system has less processing speed and greater size and weight than the other correlators. Given these limitations, the primary purpose of these investigations was to evaluate the performance of an ATR system that uses the analogue output of an optical correlator for target classification and location (correlator processing speed was considered as being of secondary importance), and factors such as system size, weight and power requirements were not addressed in this evaluation.

3.1 Description and functionality

The ATR system is a hybrid of digital electronic and optical processing hardware, and various software modules contained within a search and track program that is written in C/C++. The standard image processing algorithms are allocated to the central processing unit (CPU) of a host PC. When target recognition is required the optical correlator algorithm is implemented.

The system physically comprises the host PC, optical correlator and two image processing cards that have frame grabbing, storage and display capabilities. These are housed in the host PC and are connected to the optical correlator. Figure 2 shows a photograph of the laboratory system.




Figure 2: Photograph of laboratory concept demonstrator ATR system

Functionally the ATR system consists of an image preprocessor, correlator, postprocessor, image/filter manager and filter database. The functionality of the system, and how this is mapped onto the hardware, is illustrated in Figure 3. The functionality and hardware of each subsystem block is described in the remainder of this subsection.

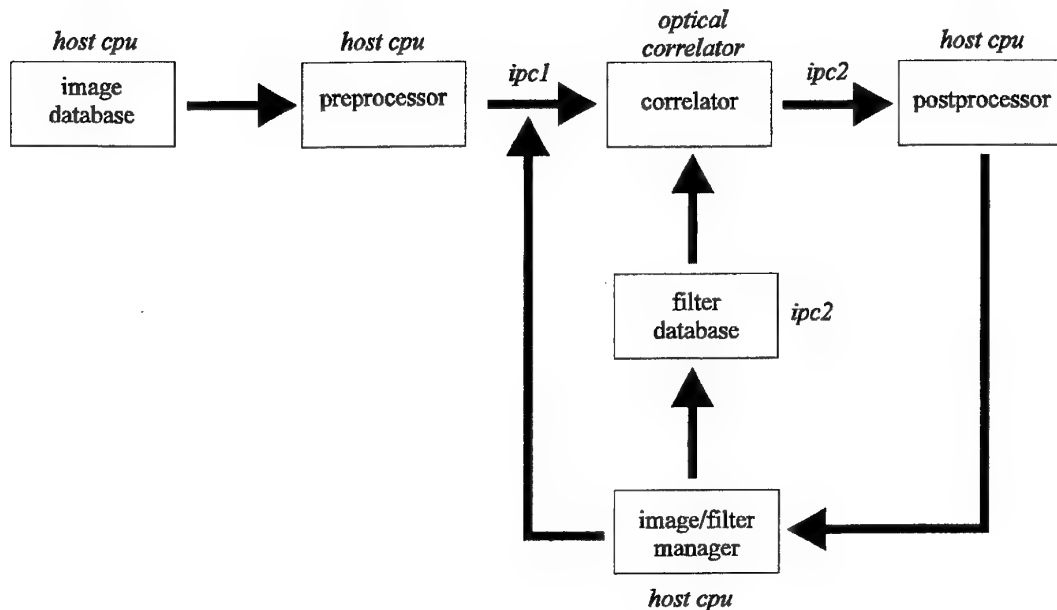


Figure 3: Block diagram showing ATR system functionality and hardware mapping (*ipcn* = image processing card *n*).

3.1.1 Image Preprocessor

The raw test imagery has to be processed by an optical correlator that will only take as input a 64×64 binary image format. Therefore, preprocessing is required to reduce the space-bandwidth product of the raw imagery, and to convert from a greyscale image to one which is binary. In addition, image artifacts present in the raw imagery have to be removed. These are due to a combination of the sensor interlacing and sensor platform motion. The sensor interlacing means that the sensor acquires one image field at a time and then interlaces two consecutive fields to generate a frame. Due to the sometimes high slew rate of the sensor on the airborne platform, movement of the sensor between field acquisitions may occur. This results in jagged target edges in the composite image frame. This effect is clearly evident in Figure 4 which shows a subimage of Figure 1 that contains the ship 2 signature. The jagged edge effect can be removed by discarding one of the fields in each frame.

The image preprocessing section takes a 512×512 8-bit pixel image and generates a 128×254 binarised edge-enhanced image. The preprocessing algorithm is implemented by the host PC and has two steps. Firstly, the space-bandwidth product of the image is reduced by vertically cropping the outer edges to give the central $390 \times$




Figure 4: Subimage of Figure 1 illustrating the jagged edge artifacts of the ship 2 signature due to sensor motion between fields.

512 section of the frame. The jagged edge effect is removed by vertically subsampling the resulting image by a factor of two to give a 390×256 image. Each row is then reduced from 390 to 130 pixels by convolving with a 3×1 moving average kernel and subsampling by a factor of three to give a 130×256 image.

In the second preprocessing step the greyscale 130×256 image is converted into an edge-enhanced binary image. The algorithm used to perform the binarisation was developed under the TOPS program and is called DIFF-3[4]. This involves scanning a 3×3 nonlinear filter over the image. For each filter position the output is the maximum minus the minimum image pixel value in the 3×3 area covered by the filter. The resulting image is called the difference image. The mean and standard deviation of the difference image are then calculated and are used to define a binarisation threshold as follows

$$Th = m + sd \times sdm \quad (1)$$

where Th is the threshold, m is the mean, sd is the standard deviation and sdm is the standard deviation multiplier. The sdm parameter is specified by the user, and so can be chosen so as to optimise the algorithm's performance. The difference image is then thresholded as follows

$$op = \begin{cases} 1 & \text{if } ip > Th \\ 0 & \text{if } ip \leq Th \end{cases} \quad (2)$$

where op is the output pixel value and ip is the corresponding pixel value in the difference image. This results in an edge-enhanced binary 128×254 image. The resulting space-bandwidth reduced image that results from the first stage of preprocessing applied to the image in Figure 1 is shown in Figure 5(a). The output from the second stage of the preprocessing is shown in Figure 5(b).




Figure 5: (a) Output of the first stage of the preprocessing. (b) Resulting edge-enhanced binary image generated by the second stage of the preprocessing.

The preprocessed image is then passed from the host PC to an image buffer in ipc1 where it is further manipulated by the image/filter manager before being input to the correlator.

3.1.2 Filter Database


The filter database is physically comprised of frame buffers in the memory of ipc2. The frame buffers store the filters needed to 'recognise' the target over the range of views likely to be encountered in the application.

The first step in filter design is to generate images containing the desired target signature views on which the filter is trained. For most applications the training target imagery will not be the same as the test imagery. There are two methods to obtain training imagery. The first method was employed by the TOPS program and involves placing the real target on a turntable [7]. A sensor, which is the same as that employed by the ATR system, is then used to gather training imagery. By rotating the turntable different azimuth views of the target can be captured and stored by the sensor. In addition to rotation of the target, the turntable can also be tilted to give


different elevation views of the target for capture and storage. Finally, the magnification of the sensor can be varied to give different target image scales. This approach is labourious, time-consuming and expensive. The cost of gathering the training image set for the TOPS program was ~ \$US 300K. In addition, it is too impracticable to put some targets, such as ships, on a turntable!

A second approach involves generating the training image synthetically by modeling the target and the sensor. This is much quicker and cheaper in comparison to the first approach, and encompasses any kind of target. The IPP discipline in LSOD is currently investigating this approach.

For this evaluation the first approach could not be employed for the reasons outlined above. Furthermore, due to the time constraints placed on the project, the second approach could not be employed either, as investigations exploring the feasibility of the approach are not yet complete. Thus, for the purposes of this evaluation it was decided to use the test imagery as the source of the training imagery. The training images were generated by firstly reducing the space-bandwidth product of the raw test image frame in the same manner as previously described in 3.1.1. The resulting image was then thresholded to remove background clutter. A 64×64 sub-image containing the target, in this case ship 2, was then extracted. Figure 6(a) shows the resulting greyscale target signature produced from Figure 5(a). To ensure the target signature in the training image was not exactly the same as the signature in the corresponding test image, the extracted subimage was then converted into a binarised edge-enhanced target signature using a Sobel filter rather than the DIFF-3 algorithm. Figure 6(b) shows the training subimage produced using the processing described above. Comparison between this and the corresponding test target signature in Figure 5(b) shows them to be slightly different - as would be the case in a real application. Using this procedure a target training image was produced for each of the 288 test image frames in the sequence.



(a)



(b)

Figure 6: (a) Greyscale target signature produced from Figure 5(a). (b) Resulting training image produced by Sobel filtering (a).

The next step is to produce the filters using the training images. For this application each filter is trained on eight contiguous target views. As there are 288 test images thirty-six filters were produced. To ensure the filters covered the range of target views in the test image sequence as illustrated in Figure 7, the first eight training images were used to train the first filter, the next eight the second filter, and so on.

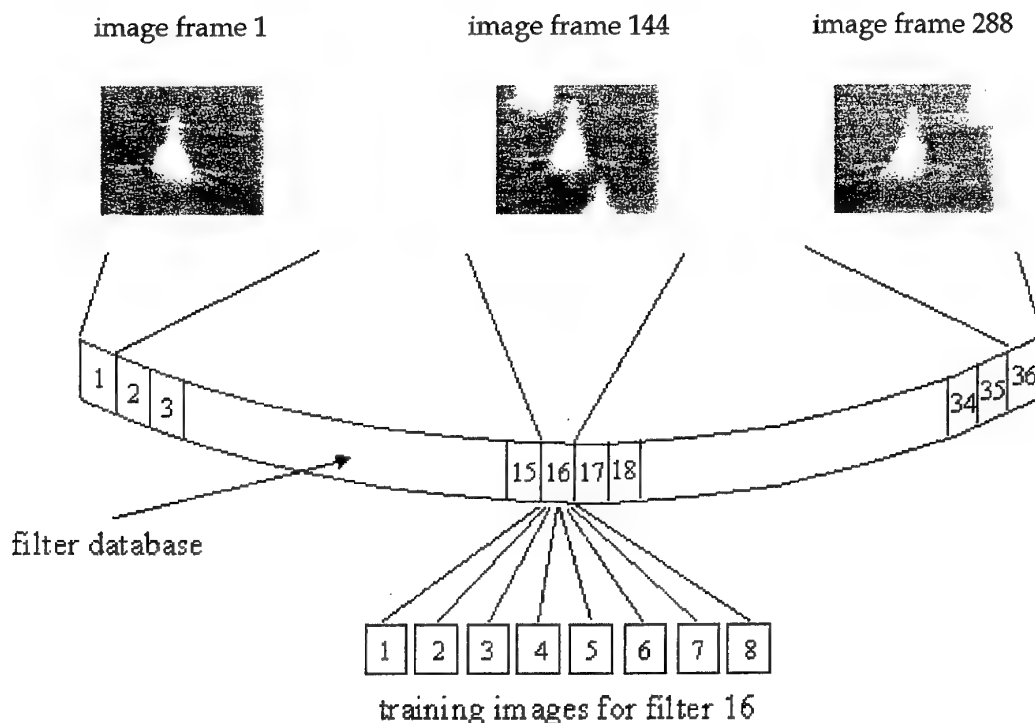


Figure 7: Diagram showing how the thirty-six filters covered the range of target views present in the test image sequence.

The training involves generating the Fourier transform of a weighted composite of each target view and binarising its phase. (The binarisation is imposed by limitations of the optical correlator hardware, but, rather fortuitously, is also desirable from a target discrimination viewpoint). A search algorithm is then used to iteratively train the weights so that the resultant filter produces the desired correlation when convolved with each of the training images. (Details of the filter design are described by Jared and Ennis [5]).

3.1.3 Correlator

The correlator section accepts a 64×64 input that is a subimage from a preprocessed test image and a 64×64 filter. The output is the convolution of the input image with the filter impulse response. As discussed previously, the correlation algorithm is

implemented by an analogue optical correlator, the details of which are described in Appendix 2. Attractive features of optical correlators include; massive parallelism, enormous processing power, small size and weight, and low power requirements.

The convolution performed by the correlator is essentially equivalent to comparing each feature in the input image with the collection of edge-enhanced reference target views in the filter impulse response. The comparison is performed by a template matching operation. Each output image from the correlator is a spatial map of the degree of likeness between the edge-enhanced reference target views and features in the input image. If a feature matches one of the edge-enhanced reference target views a strong localised correlation response, or peak, is obtained in the output image. The presence of the peak means that the feature is classified as a target. Thus, the correlator is said to have 'recognised' a target. Furthermore, the location of this peak in the output corresponds to the location of the target in the input image. If no correlation peak is obtained, then it is assumed there are no views of targets in the input image that match those of the filter impulse response. Figure 8 shows a typical test input subimage taken from the image in Figure 5(b), a filter impulse response and the correlator output.


Figure 8(c) contains a correlation peak at the position of the target in the input. The ability of the filter to produce a large correlation response is somewhat counter-intuitive given the visual nature of its impulse response. However, an explanation of this is beyond the scope of the current report. The noise component of the correlator output evident in Figure 8(c) is an inherent feature of any analogue processor, and requires the correlator output to be postprocessed.

3.1.4 Postprocessor


The correlator output is digitised and grabbed by ipc2. Following this it is passed from an image buffer in ipc2 to the host PC which processes the data to determine the presence of a correlation peak and its location. This is achieved by firstly convolving the output image with a 3×3 moving average kernel. This processing smoothes the noise component of the correlator output. Figure 9 shows the result of convolving the moving average kernel with correlator output shown in Figure 8(c).

The maximum pixel value and its location in the resulting postprocessed image is then determined. The average of a 3×3 pixel area centered on the maximum pixel value location is then calculated. This average is a measure of the correlation plane energy contained within this area and is henceforth referred to as the Energy PARAmeter (EPAR).


If the maximum pixel value is centered on a correlation peak then the EPAR value will be relatively high. Thus, if the EPAR value is above a specified threshold value then it is assumed that a correlation peak is present and so a target has been recognised. Clearly the threshold level used to decide whether a recognition occurs is another parameter whose choice can be varied to optimise system performance. Hence,



(a)



(b)



(c)

Figure 8: Test input subimage (a), filter impulse response (b), and the correlator output (c), obtained by convolving (a) with (b).




Figure 9: 3-d graphical representation of the convolution of the 3×3 moving average kernel with the correlator output shown in Figure 8(c).

provided a correlation peak is present, the postprocessor output is the EPAR value and the peak co-ordinates that correspond to the relative target position in the input image. This information is stored in a data array that is eventually passed onto the image/filter management sub-system.

3.1.5 Image/Filter Management

To acquire the target the correct filter must be matched against the target view in the current image frame. Ideally, to achieve this each preprocessed image would be correlated against the impulse response of all the filters in the database. However, as discussed, the output from the preprocessor is a 128×254 image. Unfortunately, the correlator input is only 64×64 pixels in size. Therefore, the preprocessed image needs to be partitioned in order to comply with these size restrictions. Furthermore, it would be impractical to cycle through all the filters for each correlator input as this would take too long. Thus, a search strategy has to be adopted that facilitates the 'searching' of the complete preprocessed image by as many filters as possible within the permitted time. Similarly, a track strategy has to be adopted that enables acquired targets to be tracked over continuous image frames.

The image/filter management subsystem executes the search and track strategies by controlling the input image and filter combination sent to the correlator. This subsystem is implemented by the host PC. As discussed, each preprocessed image is sent to an image buffer of size 128×256 in ipc1. During operation the system can be in either search or track mode. In search mode the first preprocessed image is subdivided by the image manager into eight sub-images as illustrated in Figure 10(a). The image manager then controls the order in which each 64×64 subimage of the image buffer is sequentially input to the optical correlator. This ensures the whole preprocessed image is covered by the search.

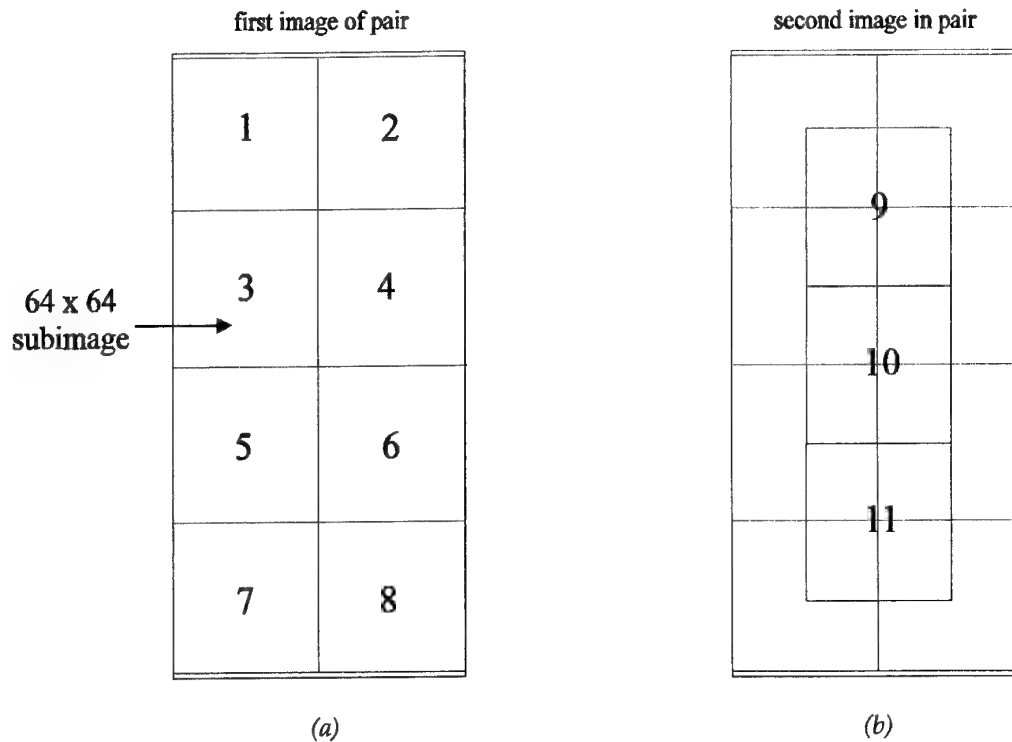


Figure 10: Diagram showing the search strategy implemented by the image manager. The strategy takes advantage of the fact that consecutive frames are similar to conduct the search over an image pair. The first image of the pair is segmented into eight subimages (a), whilst the second is segmented into three subimages (b).

For each subimage input to the correlator the filter manager sequentially sends a group of filters to the optical correlator from the filter database. As mentioned previously, the filters in the database cover a range of target views that are likely to occur in the application. (As explained in subsection 3.1.2 the filters were numbered 1 to 36 such that successively numbered filters corresponded to filters covering an adjacent set of target views). In this case it was decided to cycle through the same four filters for each subimage of an image in search mode. This was deemed to be a reasonable trade-off between the efficiency and length of time of the search.

A problem with subdividing an image into eight subimages is that the target could be positioned along one of the subimage borders. Thus, only a portion of the target would be in adjacent sub-images. This which would result in a reduced correlation response and failure to recognise the target. Based on the assumption that the target view is approximately the same in consecutive images this problem can be alleviated by subdividing the second image into three subimages as illustrated in Figure 10(b). The next four filters in the database are correlated against each subimage of the second image. This process continues whilst the system is in search mode, with consecutive disjoint pairs of frames being searched as described, and a different group of four filters being cycled through for each image frame.

In track mode only a single subimage of the image frame is passed to the correlator input by the image manager. As illustrated in Figure 11(a), this subimage is centered on the coordinates where the target was located in search mode. For each subimage in track mode only three filters are cycled through. These are the filter on which the track was initialised or maintained, and the two adjacent filters in the database. For the next image frame the subimage is centered on the coordinates in the previous subimage (Figure 11(b)). This process continues whilst the system remains in track mode.




Figure 11: Diagram illustrating tracking strategy. (a) Recognition and location of a target in search mode. (b) Resulting subimage centered on the point where the target was located in search mode, and new target location and sub-image in track mode.

3.2 Operation

The integrated system operates under the control of a search and track algorithm implemented by the host PC. In operation the system is initiated in search mode. In search mode the first image is read from the host PC hard disc and preprocessed as described in 3.1.1. The image/filter manager then implements the search strategy by segmenting the image as described in 3.1.5 and selecting a group of four filters. As there is no a priori information regarding the most likely target view the filter search was started at filter seventeen of the thirty-six. Thus, for the first image the eight sub-images were sequentially correlated with filters 17, 18, 19 and 20. Table 1 lists the filter groups cycled through for the first nine frames. After the filter group 33-36 the search wraps around to the filter group 1-4 and continues as before. The filter search is completed after nine frames and begins repeating at the tenth frame.

Table 1: Table listing the group of filters cycled through by the search strategy for each of the first nine image frames.

image frame	filter group
1	17-20
2	21-24
3	25-28
4	29-32
5	33-36
6	1-4
7	5-8
8	9-12
9	13-16

For each subimage/filter combination the correlator output was postprocessed as described in 3.1.4. The threshold value used in the postprocessing in search mode is denoted by S_{th} . If a correlation peak exceeded S_{th} , the EPAR value, peak location, subimage number and filter number were entered in an array called the target list.

The second image in the sequence is then segmented by the search strategy as discussed in 3.1.5 and the next filter group cycled through as listed in Table 1. As before the correlation outputs for each subimage/filter combination were then postprocessed. The details of any correlation peaks obtained for the second image are also entered into the target list. This leads to a possible total of $(8 + 3) \times 4 = 44$ list entries for each image pair.

If no correlation peaks are detected the system continues in search mode, subdividing alternate images and cycling through the filter database as discussed. If one or more correlation peaks are detected the maximum value of EPAR in the target list is determined. The corresponding correlation peak location and subimage number are then used to determine which of the two images the correlation peak was detected in,

and where in that image the target is. The image containing the target is then displayed on a monitor with a cross-hair placed over the target.

The system then goes into track mode and employs the track strategy described in 3.1.5. The three filters cycled through for the first image in track mode include the filter for which the maximum EPAR value was obtained in search mode, called the track filter, and the filters adjacent to this. Thus, if the maximum EPAR value in search mode was obtained with filter 23, the three filters used for the subimage in track mode would be 22, 23 and 24. As before the correlation output for each filter is postprocessed. The postprocessing threshold in track mode does not have to have the same value as in search mode, and is denoted by T_{th} . The details of any correlation peaks that exceed T_{th} are stored in the target list which is refreshed every frame.

If no correlation peaks are obtained for the first image frame in track mode the system goes back into search mode. In this case the system continues the search strategy where it left off when it was in search mode prior to going into track mode, i.e., if the last filter group in search mode was 1-4, then the system would resume the search using the filter group 5-8. If correlation peaks are obtained, the peak which gave the maximum EPAR is determined and the corresponding filter is used as the track filter for the next image frame. The correlation peak location in the current image frame is also calculated using the details stored in the target list, and this is used as the origin location for the subimage in the next image frame, as described in 3.1.5.

This process continues for as long as a correlation peak is obtained for each image frame in track mode. If, after a period, no correlation peak is obtained, the system drops back into search mode. In this case the search continues at the group of four filters which contains the track filter at which it fell out of track mode. Thus, if the track filter number was 31 then the four filters used for the first image frame back in search mode would be 29-32.

The system continues implementing this search and track algorithm until all the input image frames have been processed. During the running of the algorithm, details of the run are written to a log file which can be used afterwards for analysing the system performance. Details written to the log file include; the image frame number, the track filter number, the EPAR values obtained and their coordinates, for each frame in track mode. Also included is the number of frames spent in both search and track modes on each occasion it was in that mode.

4. System Evaluation

For the purposes of this evaluation the ATR system was tasked to acquire and track ship 2 during the test image sequence. To evaluate the system in discharging this task, the performance of the search and track algorithm was first determined with somewhat arbitrarily chosen values for the various system parameters, hereafter

referred to as the default values. Several experiments were then performed to investigate the effect of changing s_{dm} , S_{th} and T_{th} on the system's performance. Experiments were also performed in which the search strategy and the correlation peak quality parameter were changed. The results were then compared relative to those obtained with the original search and track algorithm and the default parameter values.

In addition to the evaluation experiments, a number of diagnostic tests and preliminary experiments were performed during the evaluation period. These are described in subsections 4.1, 4.2 and 4.3.

4.1 Diagnostic test

Prior to running the test sequence through the system, a diagnostic test was made in which the mean of one hundred measurements of EPAR, denoted by $EPAR_M$, was calculated. The measurements occurred continuously over a period of several minutes with the same input image and filter combination for each run. The input/filter combination used to generate the peak in Figure 8(c) was chosen as this resulted in a strong correlation peak. This test was used to provide a comparative indicator of the system performance for each run, irrespective of the parameter being investigated. Thus, provided the measured $EPAR_M$ was constant over the duration of the experiment, then differences in performance between runs could only be attributed to changes in the parameter under investigation.

4.2 Relationship between diagnostic test result variation and laboratory temperature

During the course of the evaluation experiments it was observed that the diagnostic test result varied considerably. Subjective observations indicated some correlation between the diagnostic result and the time of day, i.e., the test result seemed to be higher during the early afternoon compared to the early morning and late afternoon. As this was the warmest period of the day it was decided to determine whether this variation could be as a result of temperature fluctuations in the laboratory. Therefore, an experiment was carried out in which the diagnostic test was run every 6-7 minutes over an extended period of time, whilst simultaneously, the laboratory temperature was also measured. The latter was achieved using a device which periodically measured the temperature, digitised the result, and stored this in local memory. At the end of the specified measurement period the data in local memory could then be down-loaded to the host PC for analysis.

Six such experiments were conducted ranging in duration from 6-24 hr. The tests were usually started in the early morning just after the system had been turned on. Otherwise they started in the late afternoon when the system had been running for most of the day. Figure 12 shows the results from one of the experiments.




Figure 12: Plots of the variation in $EPAR_M$ and laboratory temperature with time.

Study of the results showed that in all cases the $EPAR_M$ was relatively low when the system was first turned on. Over the next thirty minutes the $EPAR_M$ climbed steeply to a relatively high maximum. Over the next 2-3 hr. it decreased steadily before settling to a relatively uniform level. Consequently, in order to ensure constant system performance during the evaluation experiments, they were only performed after the system had been running for a period of 24 hr. A possible reason for this start-up effect is that the analogue sections of the processing chain take some time to reach thermal stabilisation. These include the optical correlator and the video processing and analogue-to-digital converter on ipc2. A detailed explanation of the combined effect of this process for these components is beyond the scope of this investigation. In general, after stabilisation the $EPAR_M$ was observed to remain relatively constant for periods of up to six hours. Therefore, the evaluation experiments were performed over a period of 2-3 hr.

For some of the tests there was a strong correlation between temperature and the value of the $EPAR_M$, as is evident in Figure 12. Once again, this is probably due to the thermal response of the analogue components. However, for other tests the correlation with temperature was less evident. In some of these cases abrupt and smooth decreases in the $EPAR_M$ occurred coincident with an increase in temperature. A possible reason for these phenomena is the movement of components in the optical correlator due to backlash in the opto-mechanical mounts. Therefore, to ensure

optimum performance, the components of the optical correlator were re-aligned before each experiment.

Finally, although this experiment has revealed variations in the system's performance due to environmental factors, the solution to this problem was considered beyond the scope of these investigations - apart from the cautionary measures, outlined above, that were adopted to minimise these variations during the experiments,.

4.3 Filter discrimination

This preliminary experiment was performed to determine how well the correlation filters were able to discriminate between different true target views, and between true and false targets. To determine this a subimage of test image 1 that contained the target signature, shown in Figure 8(a), was correlated against each of the thirty-six filter impulse responses. The resulting correlation planes were postprocessed as described in 3.1.4, and the value of EPAR recorded. This was repeated for subimages containing signatures of ship 1 and ship 3. The variation in the correlation peak responses with the filter number is shown in Figure 13 for each ship.




Figure 13: Variation of EPAR with filter number for sub-images containing signatures of each ship.

Examination of Figure 13 reveals that the EPAR value for the true target is greater than that obtained for both false targets for all the filters. In general, the ratio of the ship 2 EPAR value to that of ship 1 or 3 is approximately in the range 2:1 to 4:3 for most filters. However, for some filters the difference between the EPAR values for the true and false targets is negligible.

An important point to note is that the EPAR value for the true target should be a maximum for filter 1, however, this is not the case. In fact, for most of the filters there is no significant difference in the EPAR value for the true target. The same can be said for both false targets.

A number of conclusions can be made from these observations. Firstly, the filter discrimination between the true and false targets is only of medium quality. Secondly, the true target may be acquired and tracked by any number of filters; not just the filter corresponding to the current target view. Finally, it was felt that too many filters were used as each filter could have been trained to cover a wider range of azimuth angles.

4.4 System performance with default parameter values

As mentioned previously, the approach adopted in the evaluation process was to first determine the system performance with the default parameter values. Due to the stochastic output component inherent in the analogue operation of an optical correlator the ATR algorithm implemented by the system need not necessarily give the same output when rerun with the same parameter values. In addition, the system is subject to changes in the laboratory environment. Therefore, to characterise the system performance with the default parameter values the test sequence was run through the system ten times over the course of the evaluation period and the results recorded.

The values of the default parameters were $s_{dm} = 1.0$, $S_{th} = 65$ and $T_{th} = 55$. At the beginning of each test run the diagnostic test described in subsection 4.1 was performed. The results from each of the runs was analysed and a number of performance indicators determined. These were: the number of frames in search mode, denoted by N_S ; the number of frames tracking a target, denoted by N_T ; and the number of frames tracking the true target and false targets, denoted by N_{TT} and N_{TF} respectively. The variation of these indicators with the diagnostic test results for each run are plotted in Figure 14.

Figure 14 shows a number of trends. Firstly, for increasing $EPAR_M$, N_S decreases and N_T increases. This agrees with what one would intuitively expect as a higher value of $EPAR_M$ should result in correspondingly higher values of EPAR during the run. This is essentially equivalent to reducing S_{th} , hence the results obtained. Secondly, N_T is significantly greater than N_S for all runs, which is rather surprising as one would not intuitively expect the search strategy employed by the system to be that efficient. Thirdly, in all cases $N_{TT} > N_{TF}$, indicating reasonable discrimination performance of




Figure 14: Variation of N_S , N_T , N_{TT} and N_{TF} with $EPAR_M$ for default parameter values.

the filters. This is consistent with the results in subsection 4.3. Finally, the overall trend for N_{TT} is a slight increase with increasing $EPAR_M$. N_{TF} remains at essentially the same level. No simple explanation can be given for these last observations as the variation of N_{TT} and N_{TF} with $EPAR_M$ has a complicated interdependence on a range of factors such as filter performance, search strategy, stochastic variation in the correlator output, S_{th} and T_{th} .

In addition to the results in Figure 14, the history of the system mode of operation for each frame was recorded, along with the track filter number for each frame in track mode. To illustrate the system history for each test run a graphical representation was employed. Figure 15 shows the graphical representation for one of the above test runs with an $EPAR_M = 83.9$.

Study of these graphical representations shows that although the track history was never exactly the same for any two runs, reflecting the nondeterministic nature of the system, trends in performance over the ten runs did emerge. Firstly, the system always acquired the true target by the 15th frame and the track filter for the initial acquisition was always in the ranges 1-5 and 30-36. Secondly, once the true target was acquired it was generally tracked to the middle of the test sequence.

However, as illustrated in Figure 15, on several occasions the system shifts from tracking the true target to tracking a false target for a small number of frames before dropping out of track mode into search mode. It would then only stay in search mode for a couple of frames before re-acquiring the true target and switching back into track mode. The short period spent in search mode before target reacquisition is a somewhat surprising result, as one would intuitively expect it to be longer. A possible




Figure 15: Graph showing a typical history of the system's mode of operation under default parameter values.

reason for this is that in some cases the true target was acquired by the 'wrong' filter, i.e., the range of target views on which the track filter was trained usually did not cover the current target view. This is consistent with the results reported in subsection 4.3. Thus, as the true target may be acquired and tracked by any number of filters, this increases the number of subimage/filter combinations that result in strong correlation responses. Consequently, the period between such combinations occurring is reduced, resulting in greater efficiency of the search strategy than one would expect. In addition, there appears to be no discernible trend of the system towards the 'correct' filter, i.e., the filter whose range of training target views matches the current target view, when the track filter is the wrong one. Instead, the system tends to remain 'stuck' in the locality of the filter with which it acquired the target. This results in the system eventually falling out of track, or jumping to a false target.

In the second half of the sequence the system falls in and out of track more frequently than in the first half, acquiring and tracking both the true target and false targets. In addition, more time is spent in search mode during this period.

During the period covering frames 150-220, the system, when in track mode, tends to track ship 1. Possible reasons for this are evident from Figure 16(a) which contains a subimage taken from frame 190. This shows that the mast of ship 2 partly obscures ship 1 leading to a merging of the two signatures. Also, the energy of ship 1 is greater than that of ship 2. These factors result in the EPAR value of the correlation peak associated with ship 1 being larger than that of ship 2 for most of this period.



(a)

(b)

Figure 16: (a) Subimage of preprocessed image (frame 190) showing merging of ship 2 with higher energy ship 1. (b) Subimage of preprocessed image (frame 250) showing similarity between end-on views of ship 2 and 3.

During the period covering frames 230-260 the system generally tracked ship 3. Figure 16(b), which is a subimage taken from frame 250, shows that the signature of ship 3 is very similar to that of ship 2 as a consequence of the view being end-on. This led to a strong correlation response for ship 3 in this period. Conversely, during this period the signature of ship 2 was generally colocated with the edge of a subimage whilst in search mode, resulting in a significantly reduced correlation response.

Finally, towards the end of each run the system always reacquires the true target and ends the run tracking the true target. It was noticeable that the system was always able to acquire and track the true target at the start and very end of the sequence. During these periods the target views are more broad side than end-on. Hence, the resulting signatures are more distinctive, which tends to facilitate easier discrimination. In the middle of the sequence the targets are end-on. As discussed above, this results in less distinctive target signatures, making discrimination difficult.

4.5 Effect of varying the sdm parameter on system performance

To investigate the effect of varying the sdm parameter three test runs were performed in which the sdm took the values zero, one and two. The runs were performed immediately after one another, with a diagnostic test being performed immediately before and after each run. The output results were recorded and analysed as described in section 4.4. All further experiments described in this section were performed, and the results analysed, in the same manner. Table 2 shows the results for these experiments.

Table 2: Table showing diagnostic tests and variation in performance indicators with sdm.


sdm	EPAR _M	N _S	N _T	N _{TT}	N _{TF}
0	90.6	66	190	13	177
1	90.9	34	249	201	48
2	89.3	28	254	232	22

From these results N_T and N_{TT} were found to increase with sdm, whilst both N_S and N_{TF} were found to decrease with increasing sdm. The variation in results was much more significant for the change in sdm from zero to one, than for one to two. The experiment was repeated with similar results being obtained.


The system operation histories with sdm equal to zero and two are shown in Figure 17. In the case of sdm equal to zero, Figure 17(a), the system tracks false targets for most of the sequence, continually falling in and out of track. In the first half of the sequence the system spends most of its time in track mode. When it falls out of track and into search mode it typically spends only two image frames in search mode before falling back into track mode. After frame 175 however, the system tends to spend more time in search mode. The inability of the system to acquire and track the true target is clearly evident.

In the case of sdm equal to two, Figure 17(b), the system tracked the true target for the first half of the sequence without falling out of track. However, it still had difficulty in maintaining track on the true target during frames 160-230. The performance in this case was significantly better than for all other cases. Furthermore, in some instances during track mode there is evidence that the true target was being tracked by the 'correct' filters over a sequence of frames.

The reasons behind these results can best be elucidated by viewing typical preprocessed input images, shown in Figure 18, obtained for each sdm value. As the sdm value increases the amount of background clutter diminishes significantly, and the fine detail edge structure of the targets becomes more resolved. Thus, for low values of sdm significant areas of the image contain clutter, leading to many false recognitions, whereas, for higher values the only features present are well resolved edge detail of the target signatures. These enhance the differences in structure, thus facilitating true recognition. For an sdm value of zero, the majority of the false targets were background clutter. In contrast, for an sdm value of one, the false targets were almost always another ship, which is consistent with the results in subsection 5.4. For an sdm value of two the false target was always another ship.



(a)



(b)

Figure 17: System operation histories with SDM equal to zero, (a), and two, (b)




Figure 18: Typical preprocessed input images obtained with sdm values of zero, (a), one, (b), and two, (c).

4.6 Influence of S_{th} and T_{th} on system performance

This experiment was performed in two stages. In the first stage S_{th} and T_{th} were progressively increased from the default values of $S_{th}/T_{th} = 65/55$ to $S_{th}/T_{th} = 75/65$ in steps of five. For each combination of S_{th} and T_{th} a test run was performed. Table 3 shows the results for these experiments.

Table 3: Table showing diagnostic tests and variation in performance indicators with values of S_{th} and T_{th} .

S_{th}/T_{th}	$EPAR_M$	N_S	N_T	N_{TT}	N_{TF}
65/55	88.8	34	242	145	97
70/60	87.5	44	232	185	47
75/65	88.6	116	155	137	18

The main outcome from the results in Table 3, is that as the values of S_{th} and T_{th} are increased N_T and N_{TF} decreases, whilst N_S increases. The value of N_{TT} increases when S_{th} is increased from sixty-five to seventy, but decreases when S_{th} is further increased.

These results are fairly intuitive, however, specific explanations can be obtained in this case by comparing the track histories. As discussed in subsection 4.4 for the default values of S_{th} and T_{th} , during the first half of the sequence the system tends to remain

'stuck' in the locality of the filter with which it acquired the target. Eventually, this results in the system directly falling out of track, or jumping to a false target for a few frames and then falling out of track. It then remains in search mode for only a few frames before re-acquiring the target. This process occurs repeatedly during the first half of the sequence.

Figure 19 shows the track history for $S_{th} = 75$ and $T_{th} = 65$. In this case, as the values of S_{th} and T_{th} increase, the higher S_{th} value tends to force the system to initially acquire the target with the correct track filter. Once again, the system tracks the target by remaining 'stuck' in the locality of the track filter. However, as T_{th} is higher the system tends to fall directly out of track sooner, without jumping onto a false target. Consequently, N_{TF} is reduced. Once back in search mode, it remains there for longer periods than was the case for lower values of S_{th} . This is because the correlation output only exceeds the higher S_{th} value when the true target is present in the current sub-image and it is matched with the correct filter. The length of time spent in search mode in this case is a truer reflection of the efficiency of the search strategy and results in an increase in N_g . Thus, each time it reacquires the target it does so in the vicinity of the correct filter. This process is repeated often during the first half of the sequence, so that each time it falls out of track the system is forced to correct itself by reacquiring




Figure 19: Graphical illustration of the track filter history with $S_{th} = 75$ and $T_{th} = 65$.

the target with the correct filters. This is illustrated in Figure 19 by the linear progression of the filter number with frame number for frames < 155 .

In the second half of the sequence, for lower values of S_{th} and T_{th} the system falls in and out of track on numerous occasions, tracking false targets for significant periods. As the value of S_{th} increases, the number of times it acquires a false target is reduced. In addition, the higher value of T_{th} ensures it only tracks a false target for a short period of time. Thus, it spends much of the second half of the sequence in search mode leading to a higher value of N_S and a reduced value of N_{TF} .

In the second stage of this experiment, the effect of varying T_{th} whilst keeping S_{th} constant was investigated. The value of S_{th} was set to 75, whilst T_{th} was set to values of 60, 65 and 70. For each combination a test run was performed. Table 4 shows the results for these experiments.

Table 4: Table showing diagnostic tests and variation in performance indicators with values of T_{th} .

S_{th}/T_{th}	$EPAR_M$	N_S	N_T	N_{TT}	N_{TF}
75/60	92.2	56	223	195	28
75/65	93.4	66	207	180	27
75/70	92.3	122	144	134	10

Study of Table 4 shows that as T_{th} increases there is a slight decrease in N_{TF} , which is low for all three cases. The value of N_S increases with increasing T_{th} . This is particularly significant when T_{th} increases from 65 to 70. In fact, there is no significant difference in the results obtained with $T_{th} = 60$, and those obtained with $T_{th} = 65$. With $S_{th} = 75$ and $T_{th} = 70$, the ratio of N_{TT} to N_{TF} obtained is the largest out of all the combinations included in Table 3 and 4. Furthermore, the value of N_{TF} is lowest for this combination. Thus, this was considered to be the optimum system performance achieved in this experiment.

Finally, comparison of the results obtained in this case for $S_{th} = 75$ and $T_{th} = 65$, with those obtained for the same combination in Table 3 shows that N_{TF} is larger and N_S is smaller in this case. These differences are qualitatively significant. A possible explanation for this is that the value obtained for the $EPAR_M$ in this case was higher, indicating higher EPAR to threshold ratios were obtained during the run.

4.7 Comparison of system performance with different search strategies

This experiment was performed to investigate the effect of using a different search strategy on the system performance. This involved the development of a new search strategy in which the same group of four filters was applied to two consecutive image frames, rather than just a single frame. The image pair was segmented into subimages as described in Figure 10. The next filter group would then be used for the ensuing image pair and so on.

To compare the two strategies a test run was performed in which the original strategy, denoted by S1, was used, followed by a second run in which the new strategy, denoted by S2, was used. In both cases $sdm = 1.0$, $S_{th} = 75$ and $T_{th} = 70$. Table 5 shows the results for these experiments.

Table 5: Table showing diagnostic tests and variation in performance indicators with different search strategies S1 and S2.

Search strategy	$EPAR_M$	N_S	N_T	N_{TT}	N_{TF}
S1	93.6	100	165	139	26
S2	93.1	92	170	161	9


From Table 5 the system performance for S2 compared with that for S1 is marginally better in relation to N_S and N_T , and significantly better in relation to N_{TT} and N_{TF} . However, these results suggest no qualitative difference in performance.

The track histories for both runs are shown in Figure 20. Comparison shows that with S2 the system started tracking on a different filter than for S1. With S2 the system tended to get 'stuck' in the locality of the same filter group for the first half of the image sequence, even though it fell in and out of track several times. In the case of S1 the system tended to move between different filter groups much more rapidly. In the second half of the sequence the system spent most of the time in search mode for both S1 and S2.


4.8 Comparison of system performance with different EPAR measurements

A factor that influences system performance is the EPAR measurement used to determine whether a correlation peak is present. In this experiment the effect of changing the EPAR measurement on system performance is investigated. For the purposes of this investigation a new EPAR measurement is employed. Rather than just determining how much energy is in a 3×3 area, the new EAPR measurement, denoted by Shape PARameter (SPAR), is a reflection of how well the shape of a correlation peak matches the following 5×5 kernel

$$SPAR = \frac{1}{25} \begin{pmatrix} -9 & -9 & -9 & -9 & -9 \\ -9 & 16 & 16 & 16 & -9 \\ -9 & 16 & 16 & 16 & -9 \\ -9 & 16 & 16 & 16 & -9 \\ -9 & -9 & -9 & -9 & -9 \end{pmatrix} \quad (3)$$



(a)



(b)

Figure 20: Track histories for S2, (a), and S1, (b).

The use of such a measurement attempts to address problems associated with false targets having more energy than the true target. For such a case the correlation peak for the true target may have less energy than that for a false target. However, the SPAR measurement will give a higher value for the true target provided its correlation peak is well-matched to the above kernel and a false one is not. The use of negative numbers means that for areas in the correlation plane with a minimum size of 5×5 pixels and a constant value, the above kernel will give a value of zero.

To investigate the effect of changing the EPAR measurement, a test run was performed in which the original EPAR was used, followed by two runs in which SPAR was used. Diagnostic tests that measured both $EPAR_M$ and $SPAR_M$ were performed immediately before and after each run. In the case of the run in which EPAR was used, the default parameter values were used, and the original search strategy, S1. For the case of the two runs in which SPAR was used, different values of S_{th} and T_{th} were employed. To determine suitable values of S_{th} and T_{th} for SPAR the default values of S_{th} and T_{th} used for EPAR were multiplied by the ratio of the $SPAR_M$ to the $EPAR_M$. Table 6 shows the results for these experiments.

Table 6: Table showing diagnostic tests and variation in performance indicators with different correlation peak measurements EPAR and SPAR.

Measurement	$EPAR_M$	$SPAR_M$	S_{th}/T_{th}	N_S	N_T	N_{TT}	N_{TF}
EPAR	93.4	15.7	65/55	6	281	229	52
SPAR	96.6	17.8	11/9	12	271	215	66
SPAR	97.6	22.2	16/14	88	169	123	46

From these results there is no qualitative difference in performance with the run employing EPAR and the first SPAR run with scaled thresholds. For these runs the value of $EPAR_M$ is very high. Consequently, very little time is spent in search mode, i.e., N_S is relatively small. In both cases, after the initial target acquisition, the system remains 'stuck' in the locality of that filter region for the sequence remainder. This is illustrated in Figure 21 which shows the track history for the first SPAR run with thresholds of 11 and 9.

For the second SPAR run in which the scaled thresholds are increased, N_S increases significantly. As a result N_T decreases, however, this is accompanied by only a small decrease in N_{TF} .

5. Synopsis of results

The main results of the last section can be succinctly stated as follows:

- Variation in the system performance indicator, $EPAR_M$, was always observed over a period of 24 hr. It was concluded that this was mainly due to fluctuations in the laboratory temperature and to movement of the optical correlator components.




Figure 21: Track history for SPAR run with thresholds 11 and 9.

- A characteristic performance of the system with default parameters was determined. The system tended to acquire and track the true target for most of the sequence. However, the track filter was often the wrong filter rather than the correct one. The system was unable to maintain track when the target became partly obscured, and was capable of better discrimination when the target views were more broad side than end-on. Tracking of false targets was unacceptably high for the default values.
- The variation of sdm was found to have a significant effect on system performance, with a low value resulting in very poor performance, and a high value resulting in the best performance. In particular, for higher values of sdm N_T and N_{TT} were high, whilst N_S and N_{TF} were very low.
- The choice of appropriate values for S_{th} and T_{th} is an important one as it has a strong influence on several performance indicators, particularly N_{TF} . Higher values of S_{th} and T_{th} resulted in an increase in the ratio N_{TT}/N_{TF} , and a decrease in N_{TF} . For these parameters there is a trade-off between minimising N_{TF} and maximising N_{TT} . In general, as one reduces N_{TF} , N_{TT} also decreases and N_S increases.

- There was no significant qualitative difference in system performance for the two search strategies, particularly in the second half of the sequence. On the other hand, although somewhat inconsistent with the results in Table 5, it was subjectively felt that S1 was the preferred strategy due to the fact that it searched through the different filter groups much more rapidly than S2.
- There was no significant difference in performance using EPAR compared with using SPAR. This indicates that the somewhat arbitrarily chosen SPAR kernel may not be that well-matched to the shape of true correlation peaks.

6. Summary and Conclusion

In this report the evaluation of an ATR system that employs an optical correlator has been described. During the evaluation the ability to acquire and track the true target over a limited range of target orientations and in the presence of low clutter and false targets was demonstrated. Furthermore, it was found that acquisition was easier if the target was somewhat broadside rather than end-on to the sensor. A shortcoming of the system was that it was not robust to minimal target obscuration. In addition, its ability to discriminate between the true and false targets was somewhat limited, and resulted in the acquisition and tracking of false targets. The discrimination performance of the system was improved by optimisation of a number of parameters, however, this also resulted in a reduction of its ability to acquire and track.

In conclusion, it is clear from this evaluation that:

- Future systems should be tested for performance variations due to environmental factors.
- Provided the preprocessing algorithm and filters perform well, this technology concept shows reasonable promise for realising automatic target acquisition and tracking in specific applications.
- Careful attention needs to be paid to filter design, particularly in relation to discrimination between targets.
- The preprocessing and thresholding on the correlation output were the most important factors investigated. The importance of a preprocessing algorithm that reduces clutter and produces distinctive well defined edge-enhanced signatures of the targets cannot be under-estimated. However, the issue of ensuring this over several operational scenarios is still unresolved. Furthermore, variation in system performance due to environmental factors has to be considered when setting thresholds.
- The use of a different search strategy and a different correlation peak quality measurement were found to have no significant difference in this case. Further

work that is beyond the scope of this evaluation needs to be performed in these areas.

It cannot be stressed too highly that the capabilities demonstrated for this surveillance scenario cannot be interpolated to others: these results are application and system specific. Furthermore, this type of system should not be seen as a panacea for real-time ATR - such a system is many years down the track and will probably be much more complex.

The application of the present type of ATR systems to different surveillance scenarios will probably require different preprocessing, collection of training images pertinent to that application, different filter design and novel filter search strategies. This will be assessed during the next evaluation stage of this project which will address the air-to-ground surveillance scenario. For this application the system will be tasked with acquiring land vehicle targets in a highly cluttered background. The test imagery will include image sequences from both visible and IR sensors that were mounted on an airborne platform. An attempt will be made to generate the target training imagery synthetically. The optical correlator will be upgraded to state-of-the-art, which should result in more efficient search strategies as a consequence of the larger input image size. Furthermore, filters will be designed by optimisation techniques implemented in-situ on the optical correlator. Finally, in addition to algorithm performance, this evaluation will address implementation issues such as processing speed.

7. References

1. Lindell, S. D. (1995) Summary of the Transfer of Optical Processing to Systems: Optical Pattern Recognition Program, *Proc. SPIE* **2489** 20-34 .
2. Liddiard, K. C. (1986) Armada 86 Surveillance Exercise - Exercise Report (Secret), *DSTO file D6332/8/1*.
3. Serati, S. A., Ewing, T. K., Serati, R. A., Johnson, K. M., and Simon, D. M. (1993) Programmable 128x128 ferroelectric-liquid-crystal spatial-light-modulator compact correlator, *Proc. SPIE* **1959**, 55-68.
4. Bauchert, K. A. (1993) Real-time hardware implementation of an optical correlation image preprocessing algorithm using an off-the-shelf image processing board, *Proc. SPIE* **1959**, 44-47.
5. Jared and D. J. Ennis (1989) Inclusion of filter modulation in synthetic-discriminant-function construction, *Appl. Opt.* **28**, 232-239.
6. Miller, P. C., and Woodruff, C., (1995) Real-time Automatic Target Recognition Using an Optical Correlator, *Proc. Electronic Technology Directions to the Year 2000*, L.C. Jain (Editor), 189-193.

7. Bauchert, K. A., Lundgreen, J. P., Flannery, D. L., and Friday, W. A. (1992) M60A2 tank turntable and test range data collection and sensitivity study for the development of smart filters, *Proc. SPIE* **1701**, 31-42.

Appendix 1

Table A.1: Performance comparison of the laboratory correlator used in these experiments with the TOPS and two state-of-the-art correlators.

	input/filter pixel size	correlation frame rate (Hz)	input/filter modulation	size (cubic cm $\times 1000$)
Laboratory correlator	64 \times 64	12.5	binary/binary	126
TOPS correlator	128 \times 128	800	binary/binary	6.5
Upgraded correlator	256 \times 256	800	binary/binary	42
OCA correlator	256 \times 256	1000	binary/binary	0.1

Table A.1 gives the performance indicators for the laboratory correlator used in these experiments, the TOPS correlator, the upgraded correlator to be used in the next phase of this task and a commercially available correlator manufactured by OCA Applied Optics in the USA. Comparing the laboratory and OCA correlators shows the image/filter size has increased by an order of magnitude, the correlation frame rate has increased by two orders of magnitude and the physical size of the correlators has been reduced by three orders of magnitude. Further improvements are expected as devices have recently become commercially available that will allow greyscale correlators to be developed with frame rates of 10 KHz. The OCA correlator is mounted on a conventional printed circuit board that fits inside a PC. This device has only recently become commercially available at an initial price of \$US 100 K. The company expects a mature price of around \$US 5 K. The upgraded correlator developed for the next phase of the task uses the same components and technology as the OCA correlator.

Appendix 2

A schematic diagram of the optical correlator used in this evaluation is shown in Figure A.1. The correlator consists of a low-power Helium-Neon (He-Ne) laser, an optical system, two liquid crystal televisions (LCTVs) and a CCD.

The LCTVs are miniature real-time devices that display video imagery. For this application they have been modified to operate in certain modes. LCTV1 operates in a binary amplitude mode, i.e. the LCTV pixels are either transparent or opaque to the laser light, whilst LCTV2 operates in a binary-phase mode, i.e. the LCTV pixels are transparent but either retard or advance the phase of the laser light. Input to LCTV1 and LCTV2 is provided by ipc1 and ipc2 respectively.

The operation of the correlator is as follows. The beam from the He-Ne laser is expanded and passes through L1 to illuminate LCTV1 in the input plane of the correlator. Following transmission through LCTV1 the binarised edge-enhanced image displayed on LCTV1 is written onto the laser beam. The telephoto lens arrangement consisting of L1 and L2 produces an optical Fourier transform (FT) of the input image displayed on LCTV1. It is this ability of simple lenses to produce the




Figure A.1: Schematic diagram of the optical correlator used in the evaluation.

optical FT of a coherent image that is responsible for the processing power of analogue optical processors.

LCTV2 is positioned in the same plane, called the filter plane, as the optical FT. The telephoto lens arrangement ensures the scale of the optical FT is matched to the scale of the filter displayed on LCTV2, whilst minimising the length of the optical train. On transmission through LCTV2 the optical FT is spatially phase modulated according to the pattern described by the binary filter displayed on LCTV2. Another FT lens, L3, positioned after LCTV2 performs a second FT to produce a filtered image which is detected by a CCD positioned in the correlator output plane.

With an image and filter displayed on LCTV1 and 2 respectively, the output of the correlator is the convolution of the input image with the impulse response of the filter. If the filter is derived from a matched filter the output is a correlation. The correlation is detected by the CCD whose video output is grabbed by ipc2 and transferred to the host PC for postprocessing.

Evaluation of a Laboratory Automatic Target Recognition System: Air-to-Sea Image Sequence

Paul Miller, Peter Virgo, Mike Royce and Steve Angeli

AUSTRALIA

1. DEFENCE ORGANISATION

a. Director Combat Force Development (Land)

b. S&T Program

Chief Defence Scientist	}
FAS Science Policy	} shared copy
AS Science, Industry and External Affairs	
AS Science Corporate Management	}
Counsellor Defence Science, London (Doc Data Sheet)	
Counsellor Defence Science, Washington (Doc Data Sheet)	
Scientific Adviser to Thailand MRDC (Doc Data Sheet)	
Director General Scientific Advisers and Trials/Scientific Adviser Policy and Command (shared copy)	
Navy Scientific Adviser	
Scientific Adviser - Army	
Air Force Scientific Adviser	
Director Trials	

Electronics and Surveillance Research Laboratory

Director

Chief Land, Space and Optoelectronics Division

Research Leader Land Systems (LSOD)

Head of Image Processing and Propagation (LSOD)

Head of System Test and Evaluation (LSOD)

Head Advanced System Concepts (LSOD)

Head Land Operations Analysis (LSOD)

Dr D. Bertilone (LSOD)

Dr A. Yakovleff (LSOD)

M. Fiebig (LSOD)

Author(s): Paul Miller, Peter Virgo, Mike Royce and Steve Angeli

Head Radar Land Surveillance (MRD)

Aeronautical and Maritime Research Laboratory

Director

Head Electro-optic Seekers, Weapons Systems Division

DSTO Library

Library Fishermens Bend

Library Maribyrnong

Library DSTOS (2 copies)

Library, MOD, Pyrmont (Doc Data sheet)

c. Forces Executive

Director General Force Development (Sea),

SOEW in FD (Sea),

Director General Force Development (Land),

Director General Force Development (Air),

Director General Force Development (Joint)

d. Navy

DAVPROJ-N TECH2

e. Army

ABCA Office, G-1-34, Russell Offices, Canberra (4 copies)

SO (Science), HQ 1 Division, Milpo, Enoggera, Qld 4057 (Doc Data Sheet)

NAPOC QWG Engineer NBCD c/- DENGERS-A, HQ Engineer Centre Liverpool Military Area, NSW 2174 (Doc Data Sheet)

DARMD-A

f. Air Force

g. S&I Program

Defence Intelligence Organisation

Library, Defence Signals Directorate (Doc Data Sheet)

h. Acquisition and Logistics Program

i. B&M Program (libraries)

OIC TRS, Defence Central Library

Officer in Charge, Document Exchange Centre (DEC), 1 copy

DEC requires the following copies of public release reports to meet exchange agreements under their management:

*US Defence Technical Information Centre, 2 copies

*UK Defence Research Information Centre, 2 copies

*Canada Defence Scientific Information Service, 1 copy

*NZ Defence Information Centre, 1 copy

National Library of Australia, 1 copy

2. UNIVERSITIES AND COLLEGES

Australian Defence Force Academy

Library

Head of Aerospace and Mechanical Engineering

Deakin University, Serials Section (M list)), Deakin University Library,
Geelong, 3217,

Senior Librarian, Hargrave Library, Monash University

Librarian, Flinders University

3. OTHER ORGANISATIONS

NASA (Canberra)

AGPS

State Library of South Australia

Parliamentary Library, South Australia

OUTSIDE AUSTRALIA

4. ABSTRACTING AND INFORMATION ORGANISATIONS

INSPEC: Acquisitions Section Institution of Electrical Engineers

Library, Chemical Abstracts Reference Service

Engineering Societies Library, US

Materials Information, Cambridge Scientific Abstracts

Documents Librarian, The Center for Research Libraries, US

5. INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK

Library - Exchange Desk, National Institute of Standards and
Technology, US

SPARES (7 copies)

Total number of copies: 72

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA					
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE Evaluation of a Laboratory Automatic Target Recognition system: Air-to-Sea Image Sequence			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Paul Miller, Peter Virgo, Mike Royce and Steve Angeli			5. CORPORATE AUTHOR Electronics and Surveillance Research Laboratory PO Box 1500 Salisbury SA 5108		
6a. DSTO NUMBER DSTO-TR-0478		6b. AR NUMBER AR-010-106		6c. TYPE OF REPORT Technical Report	
7. DOCUMENT DATE February 1997					
8. FILE NUMBER D 9505-10-123		9. TASK NUMBER ADL 94/128		10. TASK SPONSOR DCFD(L)	
				11. NO. OF PAGES 38	
				12. NO. OF REFERENCES 7	
13. DOWNGRADING/DELIMITING INSTRUCTIONS				14. RELEASE AUTHORITY Chief, Land Space and Optoelectronics Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DIS NETWORK OFFICE, DEPT OF DEFENCE, CAMPBELL PARK OFFICES, CANBERRA ACT 2600					
16. DELIBERATE ANNOUNCEMENT No limitations					
17. CASUAL ANNOUNCEMENT Yes					
18. DEFTEST DESCRIPTORS Automatic target recognition Optical correlators Image Processing					
19. ABSTRACT The operation and evaluation of a laboratory automatic target recognition (ATR) system is described. The ATR system was evaluated against its ability to acquire and track a specific target in the presence of false targets and background clutter. A ten second air-to-sea image sequence that contained the infrared signatures of three ship targets and sea clutter was used as the test sequence. The evaluation involved assessing the system performance with default values for various system parameters. These were then varied and the effect on system performance measured. The ability to acquire and track the true target over a limited range of target orientations and in the presence of low clutter and false targets was demonstrated. Furthermore, it was found that acquisition was easier if the target was somewhat broadside rather than end-on to the sensor. A shortcoming of the system was that it was not robust to minimal target obscuration. In addition, its ability to discriminate between the true and false targets was somewhat limited, and resulted in the acquisition and tracking of false targets. The evaluation showed that the system discrimination was improved by optimisation of a number of parameters, however, this also resulted in a reduction of its ability to acquire and track the true target. In addition, those parameters that were most important in terms of their effect on the system performance were determined.					